The Fermilab HEPCloud Facility

Gabriele Garzoglio, on behalf of the HEPCloud Leadership Team:

Stu Fuess, Burt Holzman, Rob Kennedy,
Steve Timm, Tony Tiradani

IIT Workshop

30 Mar 2016

Computing at the Fermi National Accelerator Laboratory

- Lead United States particle physics laboratory
  - ~100 PB of data on tape
  - High Throughput Computing characterized by:
    - “Pleasingly parallel” tasks
    - High CPU instruction / Bytes IO ratio
    - But still lots of I/O. See Pfister: “In Search of Clusters”
- Focus on Neutrino Physics
  - Including the NOvA Experiment
- Strong collaborations with international laboratories
  - CERN / Large Hardron Collider (LHC) Experiments
  - Lead institution for the Central Muon Solenoid (CMS)

Why HEPCloud?

Computing Demand for Science Keeps Growing

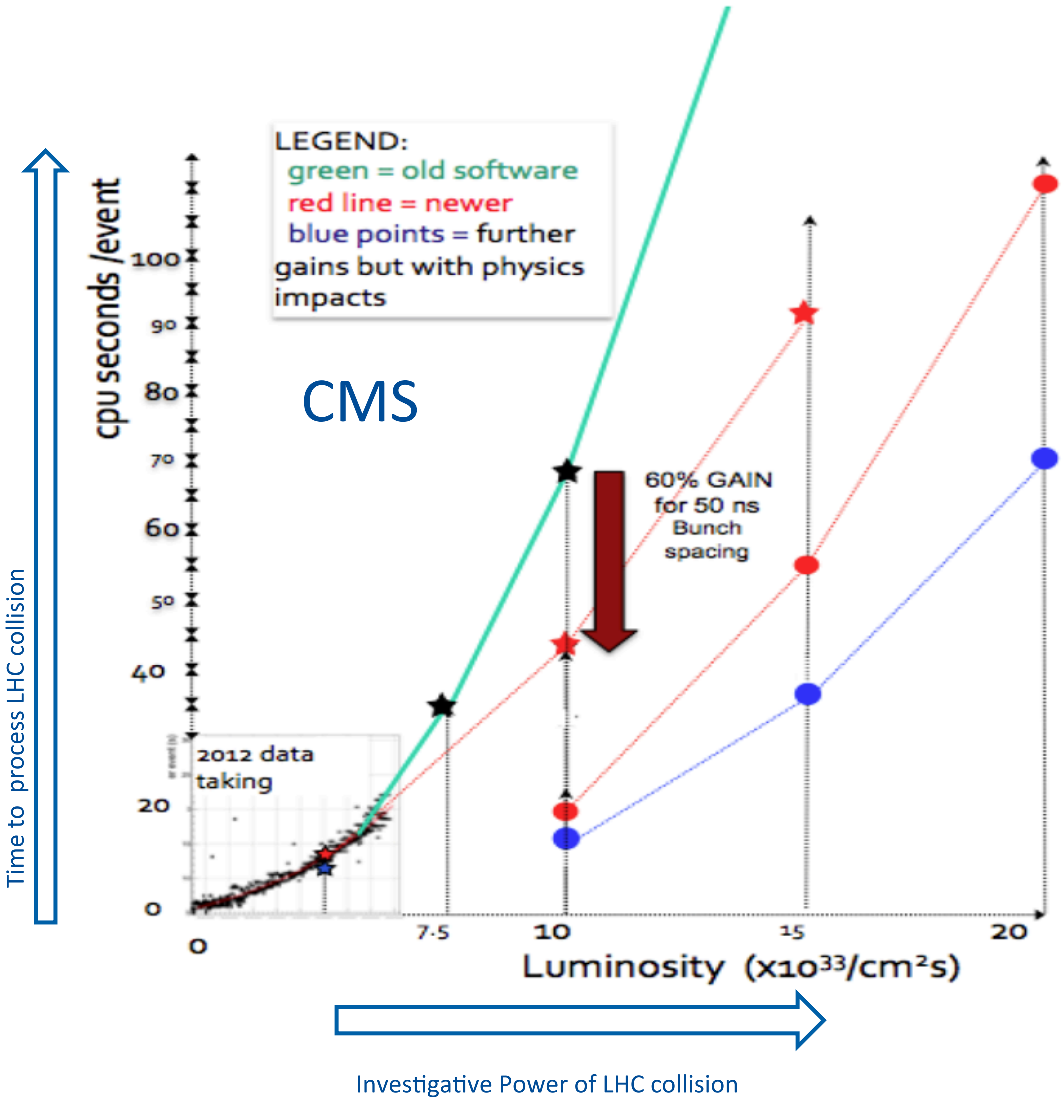

- Resource demands for science will increase significantly while instruments get more powerful and probe nature in more detail
  - Development is making big strides in improving the software performance
  - But in the end, the problem to provision resources to fulfill demand is a daunting task ➜ and it can only get worse!
- Question: How can we provide access to sufficient resources to do science in the future? ➜ **CAPACITY**

Use of Resources is Cyclic

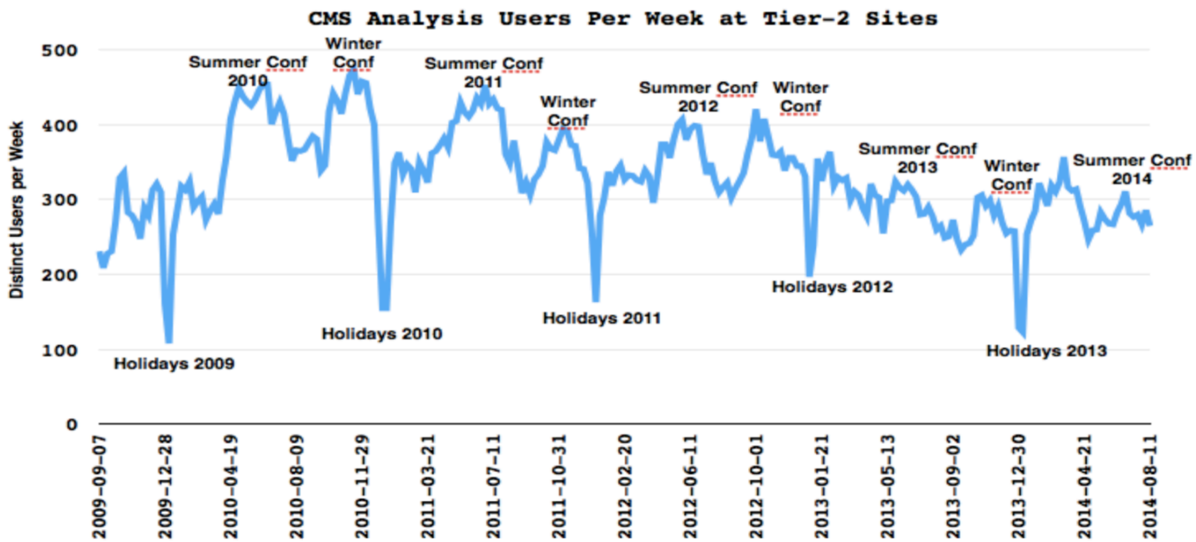

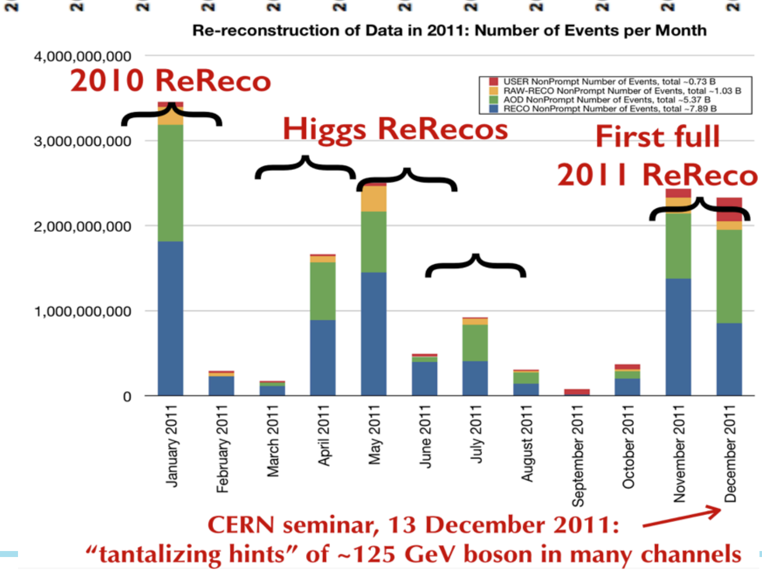

- The activity of the experiments is not constant!
  - It varies significantly with external triggers
    - Instrument operation schedule
    - Conference schedule
    - Holiday festivities, vacation time, etc.

- Question: How can we provision resources efficiently? ➜ **Elasticity**
  - Provide efficiently resources if there is demand
  - Don't waste resources if there isn't.

Moving beyond "Provisioning for Average"

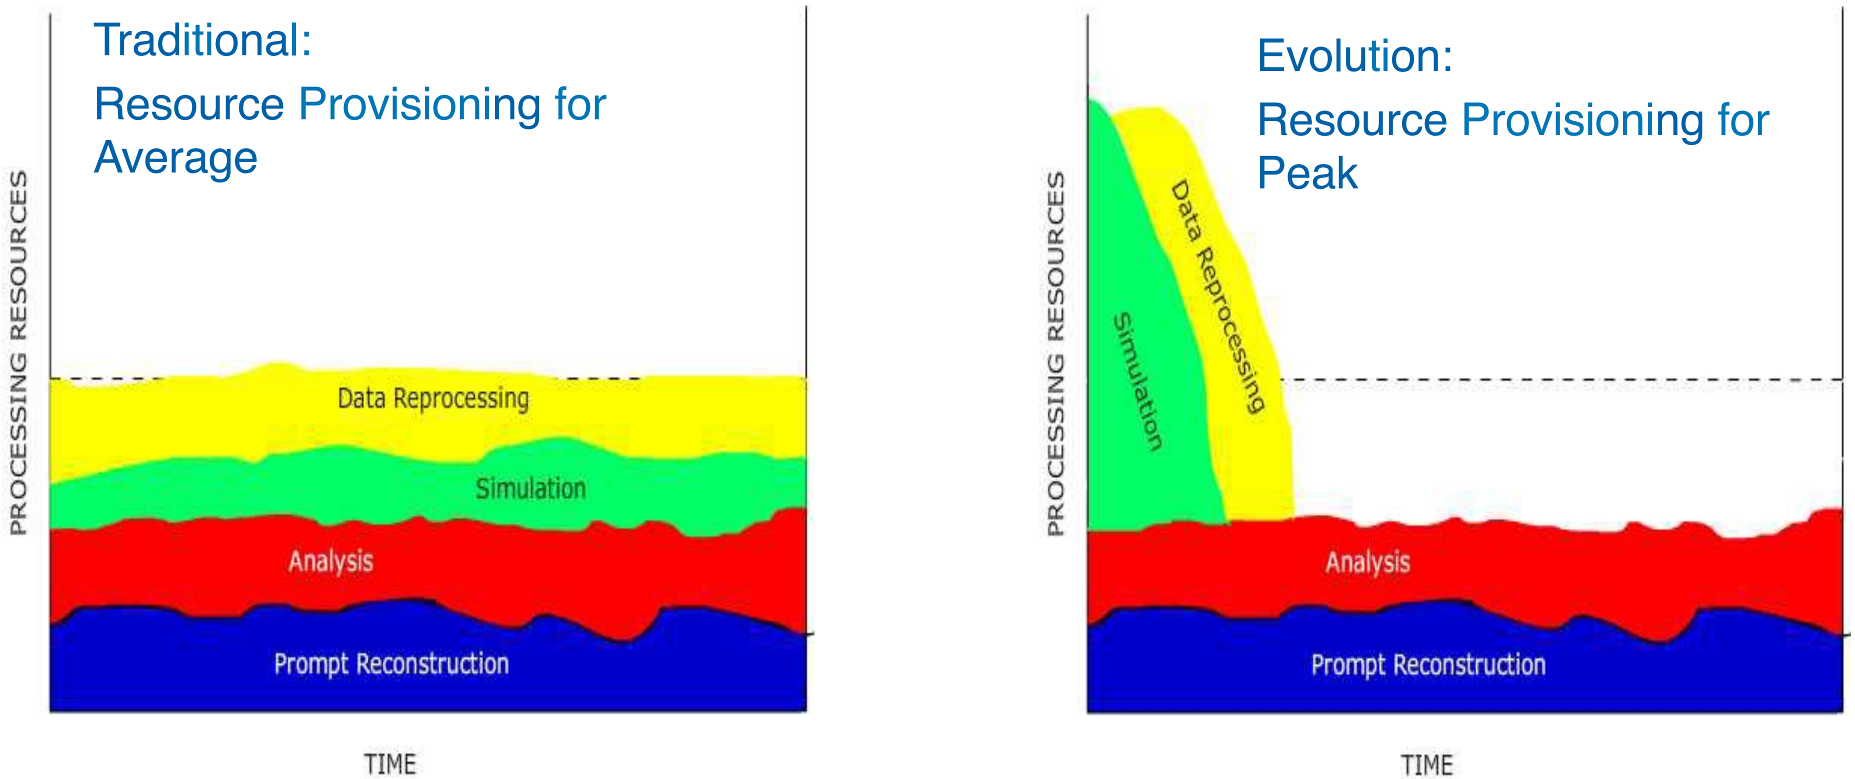

- Experiments don't need all the resources all the time
    - provisioning needs to be adaptable, providing facility "elasticity" to handle "burst usage"

  ➔ ***incorporate and manage*** "rental" resources

Classes of Resource Providers

Grid

- Virtual Organizations (VOs) of users trusted by Grid sites
- VOs get allocations ➜ **Pledges**
  - Unused allocations: opportunistic resources

"Things you borrow"

Trust Federation

Cloud

- Community Clouds - Similar trust federation to Grids
- Commercial Clouds - **Pay-As-You-Go** model
  - Strongly accounted
  - Near-infinite capacity ➜ **Elasticity**
  - Spot price market

"Things you rent"

Economic Model

HPC

- Researchers granted access to HPC installations
- Peer review committees award **Allocations**
  - Awards model designed for individual PIs rather than large collaborations

"Things you are given"

Grant Allocation

What is HEPCloud?

HEP Cloud Architecture

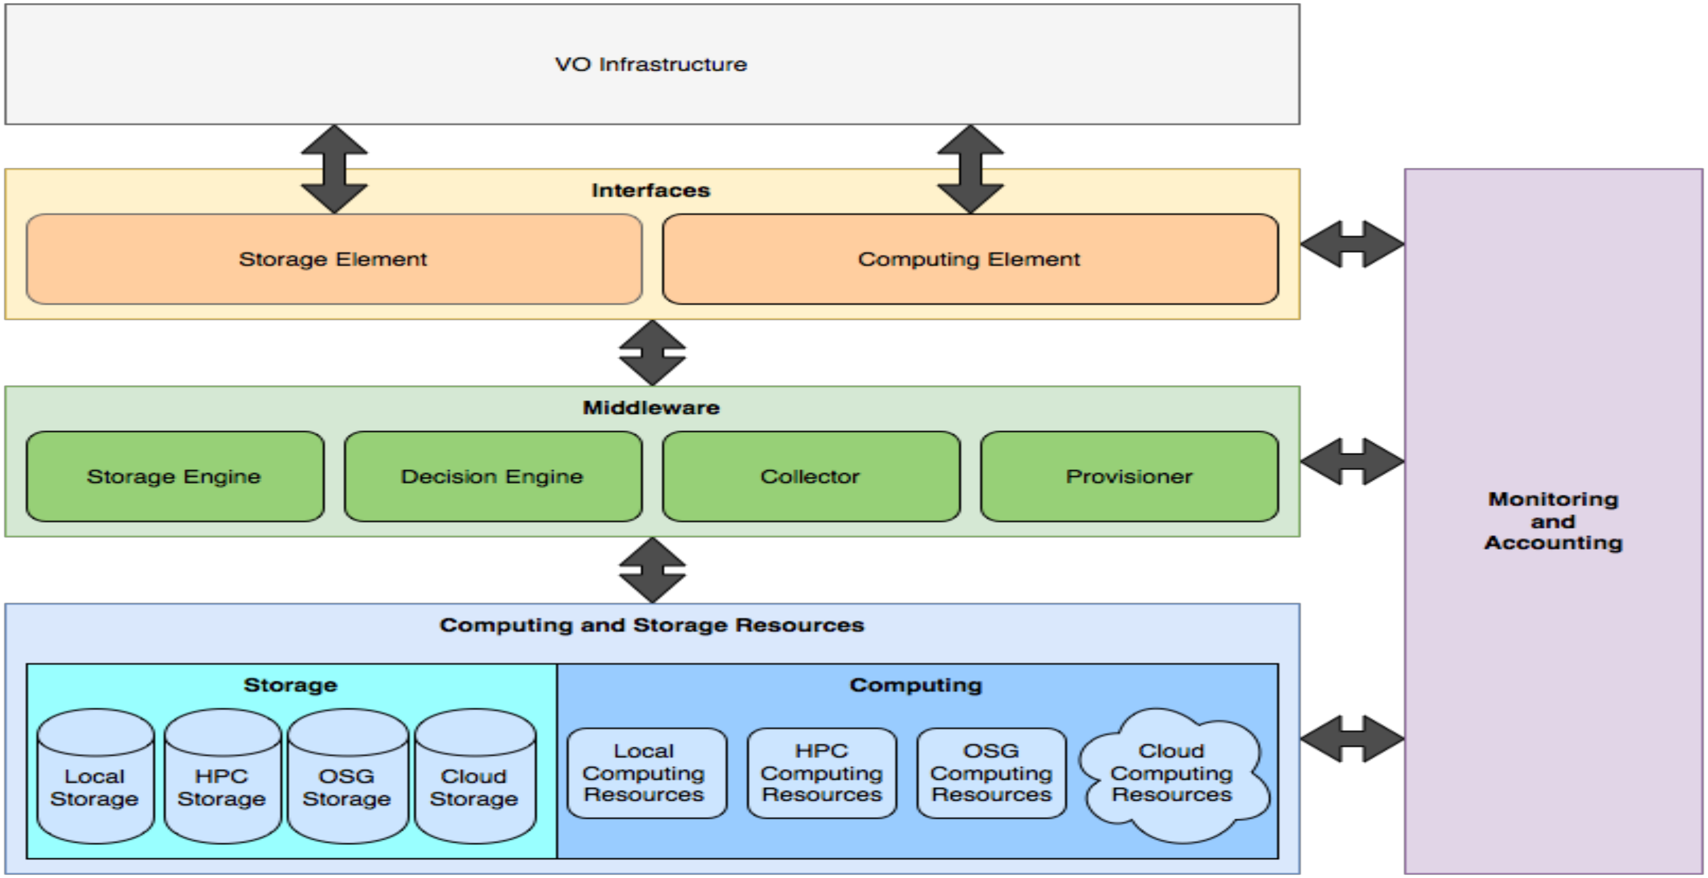

- Provision commercial cloud resources in addition to physically owned resources
- Transparent to the user
- Pilot project / R&D phase

Fermilab

Fermilab HEPCloud: expanding to the Cloud

- Where to start?
  - Market leader: Amazon Web Services (AWS)

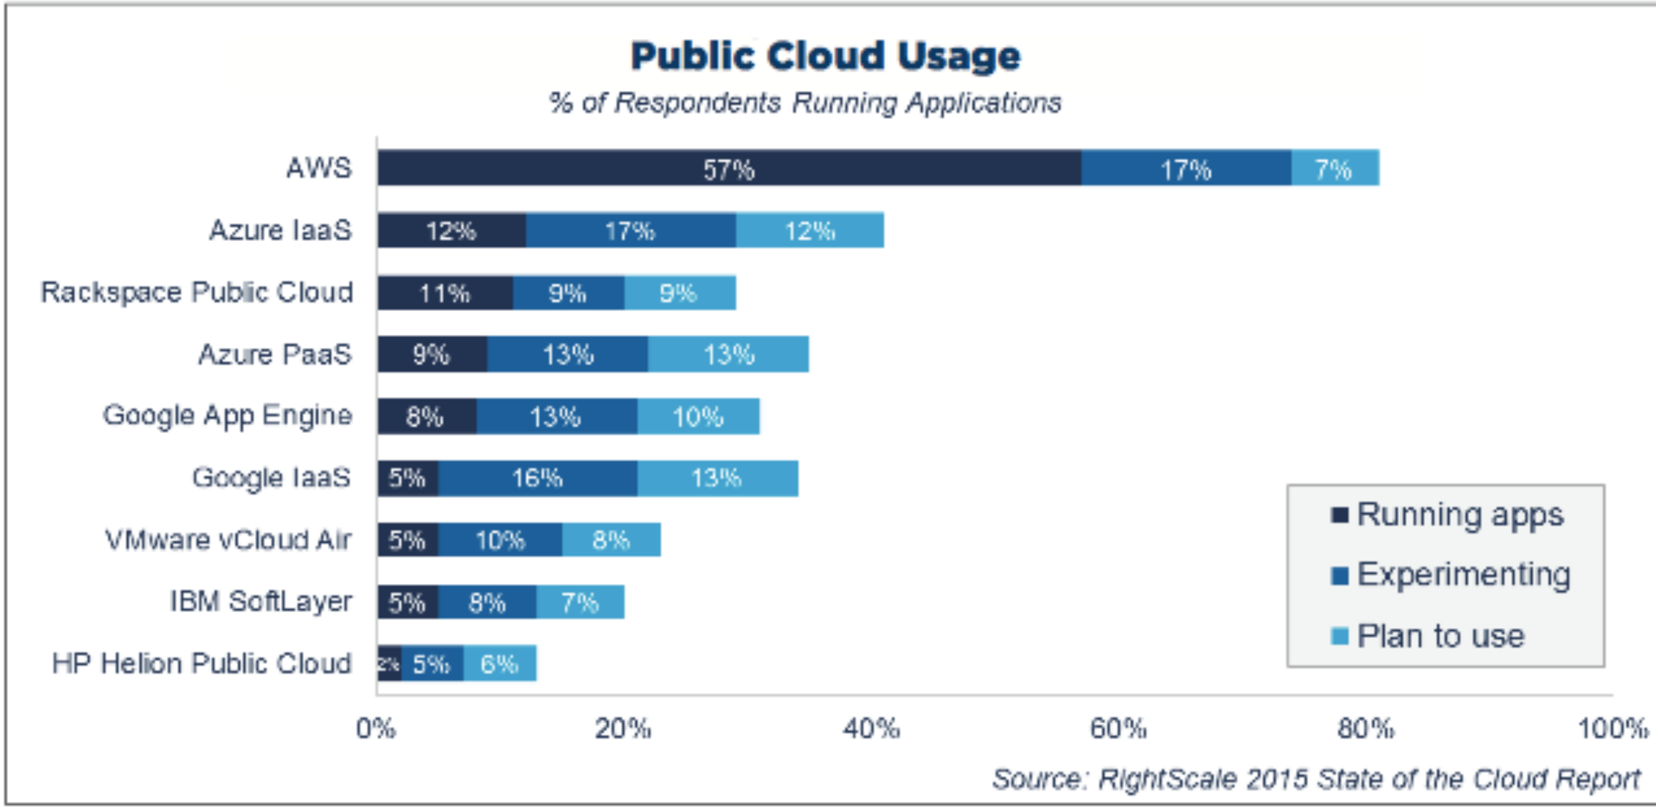

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

Integration challenges

- What are the challenges we face to run at scale on AWS?
- Many HEP experiments and University / Lab staff have been working with AWS to improve the integration of computing resources to run scientific workflows
  - Examples: Atlas, CMS, STAR, NOvA, etc. / BNL, FNAL, etc.
  - It's all about understanding how most efficiently to use AWS capabilities
- Several areas of work
  - Performance
  - Networking
  - Provisioning
  - Storage
  - On-demand Services
  - Monitoring and Accounting
  - Cost containment

Integration Challenges: Performance

- Benchmarks are used to compare workflow duration on AWS (and $$) with local execution

Supported by Davide Grassano (INFN)

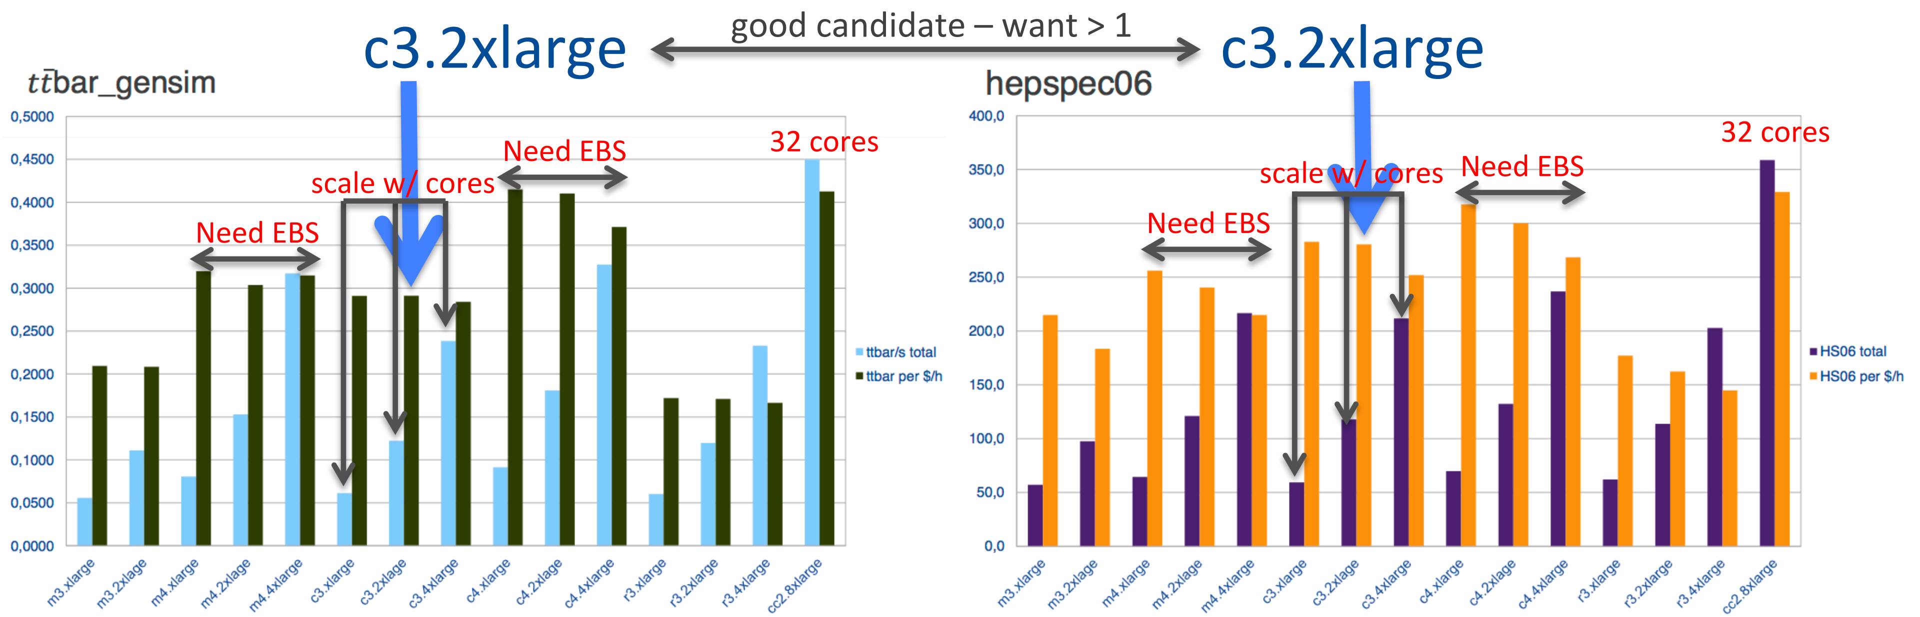

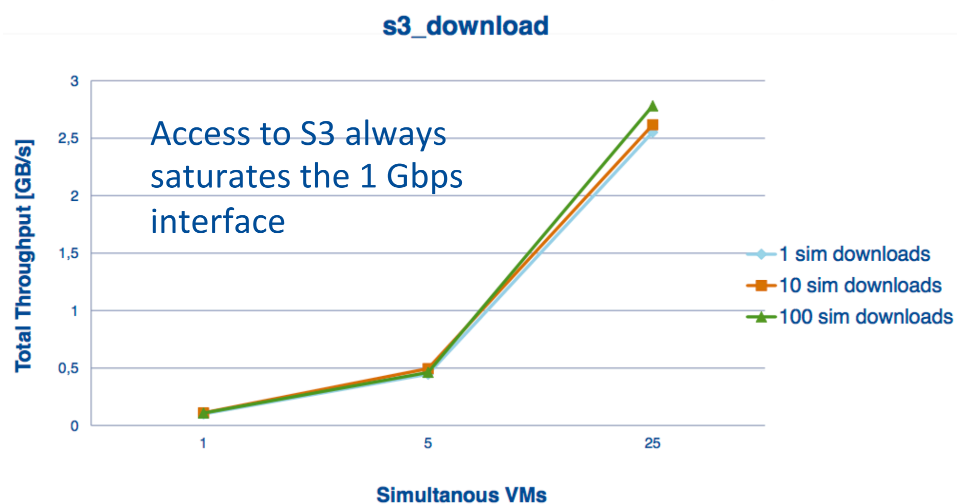

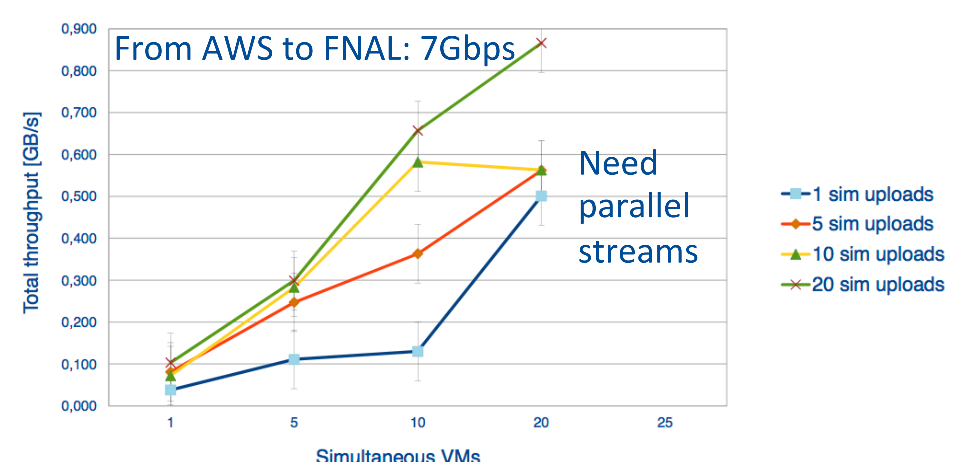

Integration Challenges: Provisioning

- AWS has a fixed price per hour (rates vary by machine type)
- Excess capacity is released to the free (“spot”) market at a fraction of the on-demand price
  - End user chooses a bid price and pays the market price. If price too high → eviction
- The Decision Engine oversees the costs and optimizing VM placement using the status of the facility, the historical prices, and the job characteristics.

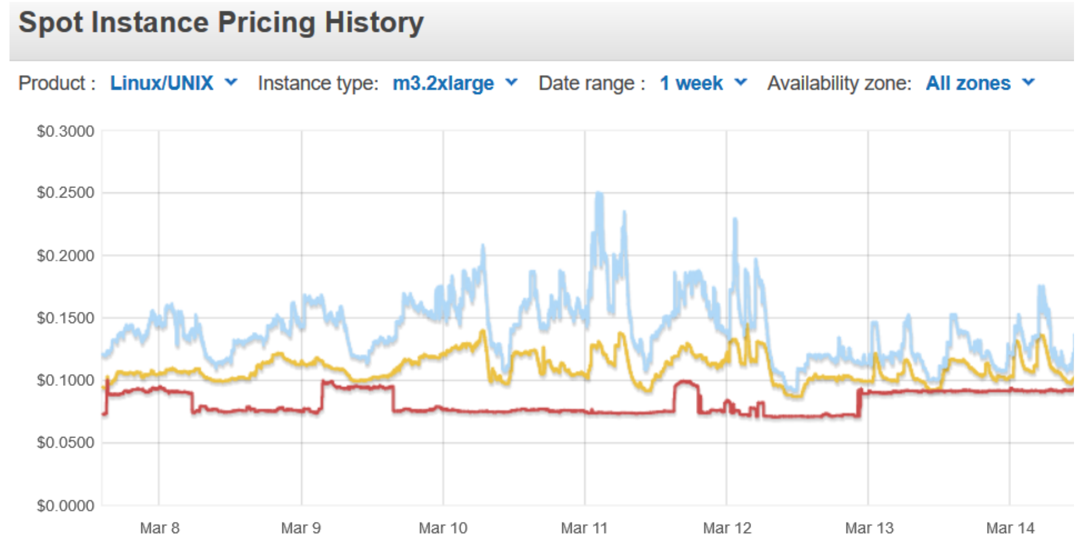

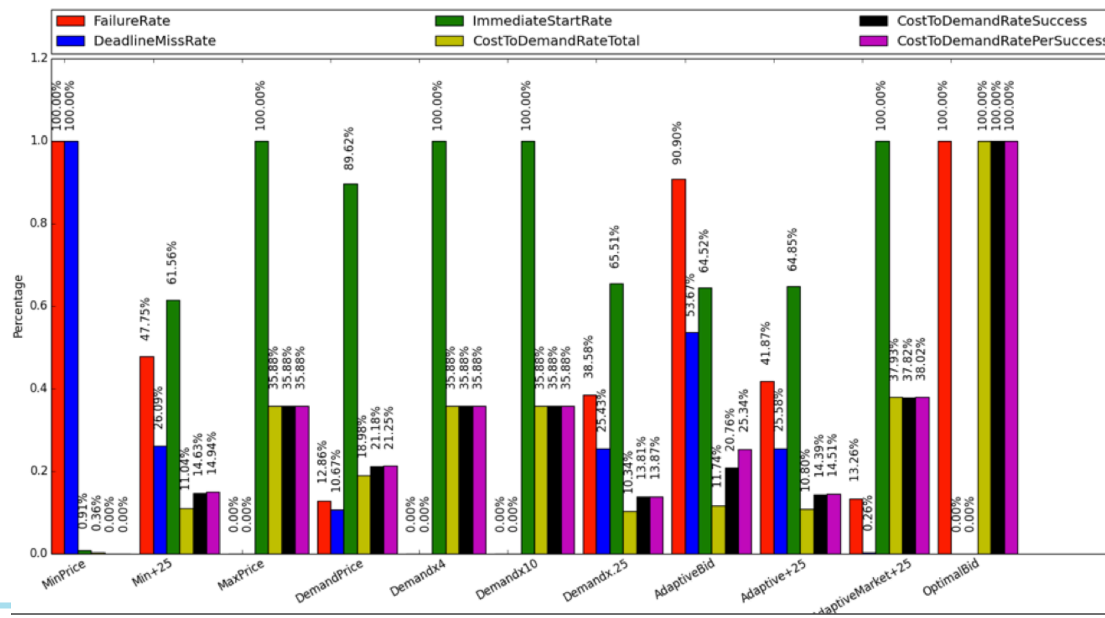

Hao Wu (IIT) PhD dissertation

Integration Challenges: Image Portability

- Build "Golden Image" from standard Fermilab Worker Node configuration.
- Build VM management tool, considering:
  - HVM virtualization (HW VM + Xen) on AWS: gives access to all AWS resources
  - Contain VM size (saves import time and cost)
  - Import process covers multiple AWS accounts and regions
  - AuthN with AWS use short-lived role-based tokens, rather than long term keys

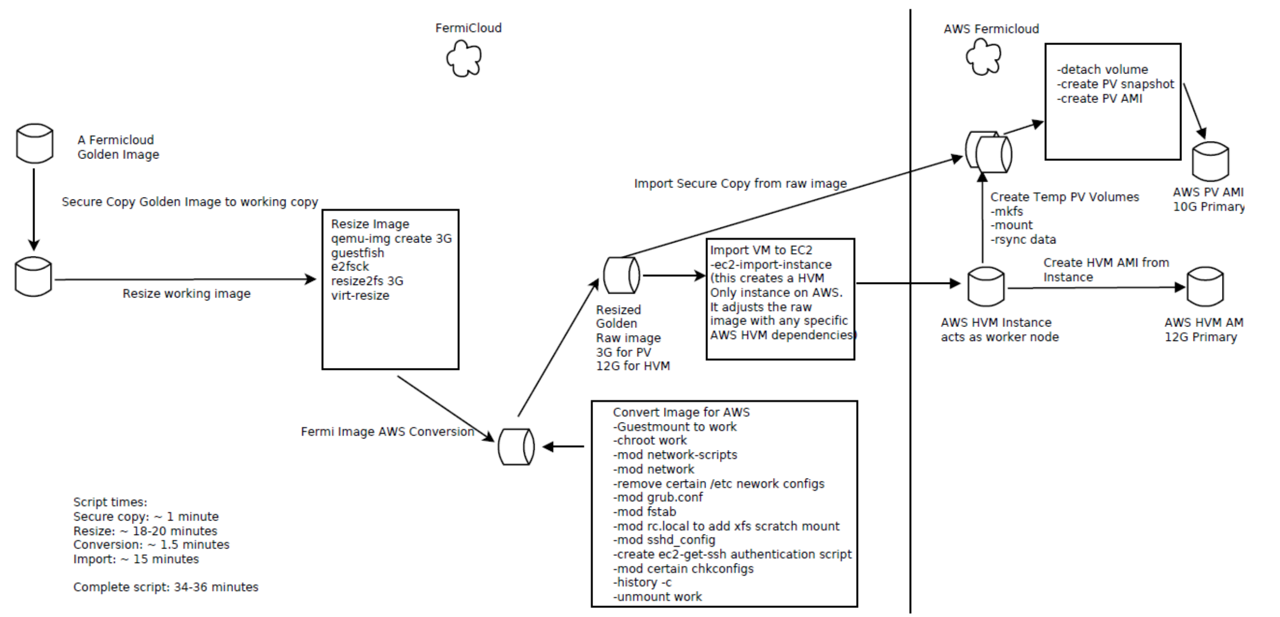

Integration Challenges: On-demand Services

- Jobs depend on software services to run
- Automating the deployment of these services on AWS on-demand - enables scalability and cost savings
  - Services include data caching (e.g. Squid) WMS , submission service, data transfer, etc.
  - As services are made deployable on-demand, instantiate ensemble of services together (e.g. through AWS CloudFormation)
- Example: on-demand Squid
  - Deploy Squid via auto-scaling services. Squid is deployed if average group bandwidth utilization is too high. Server is deployed or destroyed in 30 seconds.
  - Front Squids with an Elastic Load Balancer (ELB).
  - Name the ELB for that region via Rt. 53

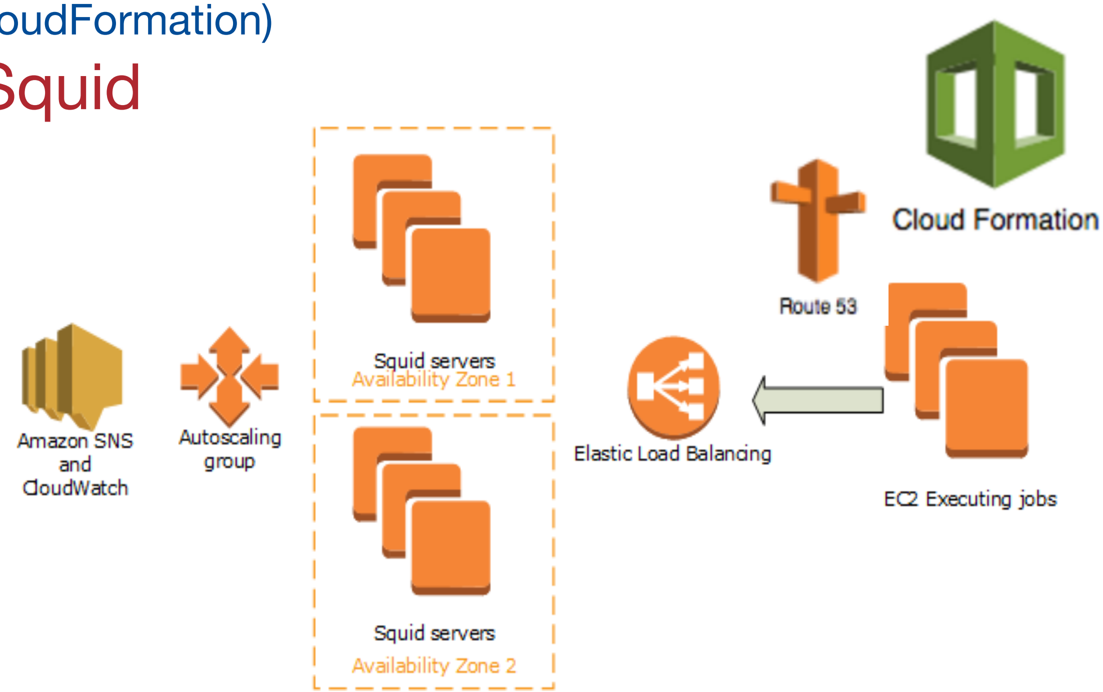

Supported by Sandeep Palur (IIT), Rahul Krishnamurthy (IIT) and Claudio Pontili (CNI)

Integration Challenges: Networking

- Implement routing /firewall configuration to utilize peered ESNet / AWS to route data flow through ESnet
- AWS / ESNet data egress cost waiver
  - For data transferred through ESNet, transfer charges are waived for data costs up to 15% of the total
- Topology: 3 AWS Regions in the US
  - Each region with multiple Availability zones and instance types

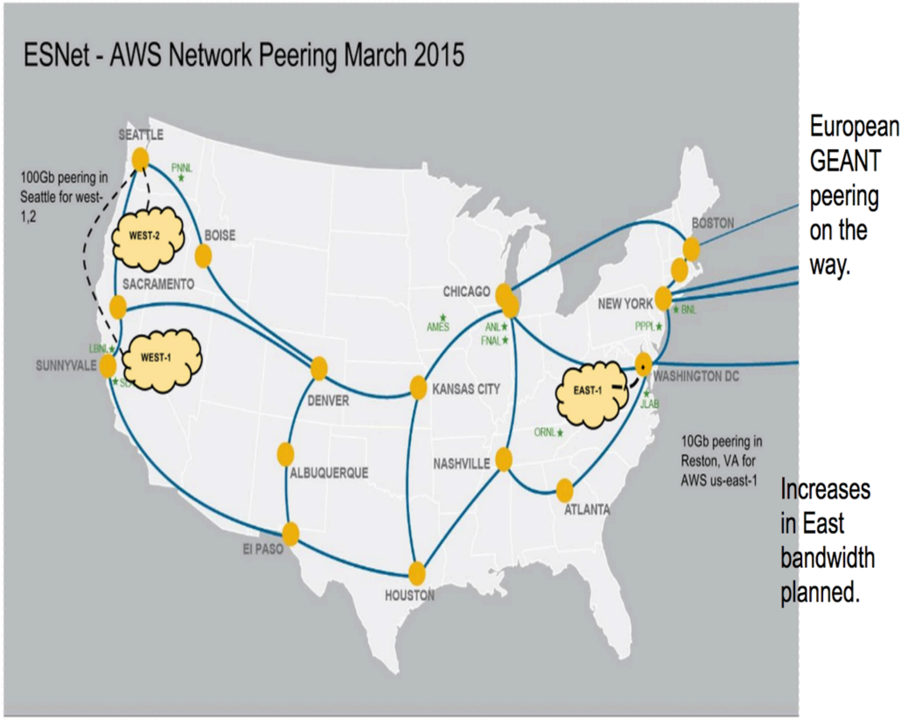

Integration Challenges: Storage

- Integrate S3 storage stage-in/-out for AWS internal / external access - enables flexibility on data management
  - Consider O(1000) jobs finishing on the Cloud and transferring output to remote storage
  - Storage bandwidth capacity is limited
  - 2 main strategies for data transfers
    - 1) Fill the available network transfer by having some jobs wait - Put the jobs on a queue and transfer data from as many jobs as possible - idle VMs have a cost
    - 2) Store data on S3 almost concurrently (due to high scalability) and transfer data back asynchronously - data on S3 have a cost
  - The cheapest strategy depends on the storage bandwidth, number of jobs, etc.

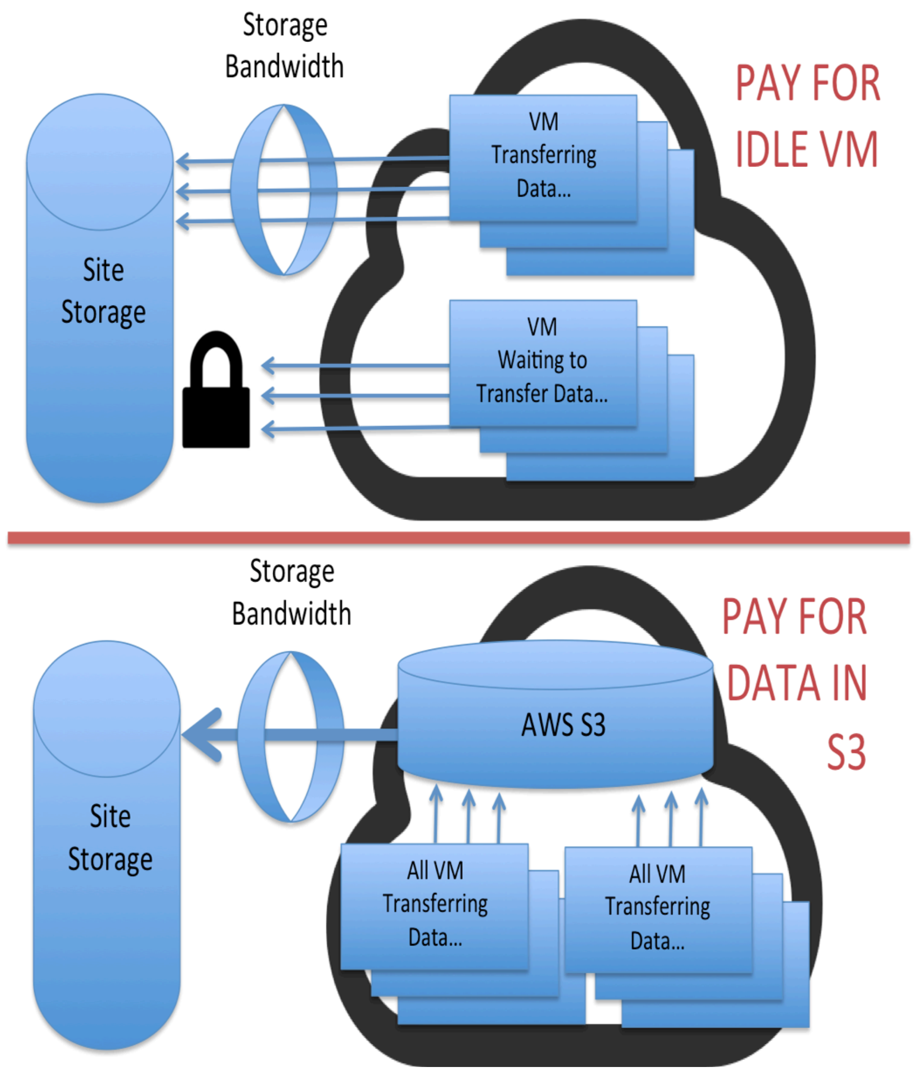

Integration Challenges: Monitoring and Accounting

Monitor # GCloud VMs

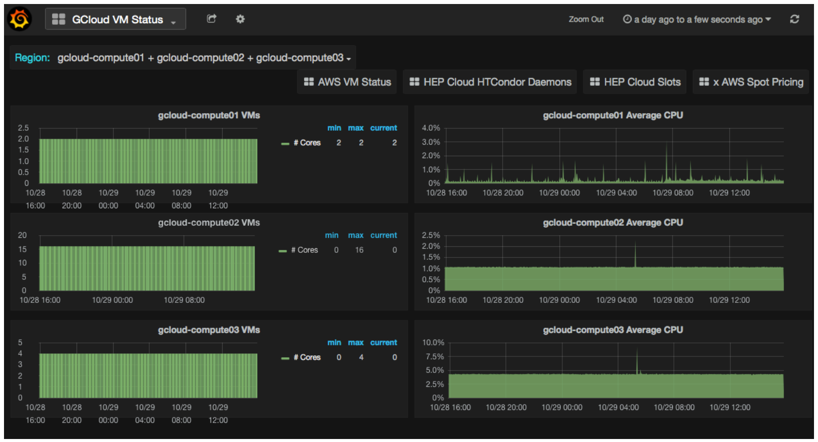

Monitor HEP Cloud Slots

Monitor # AWS VMs

Accounting:
$ by VO and VM Type

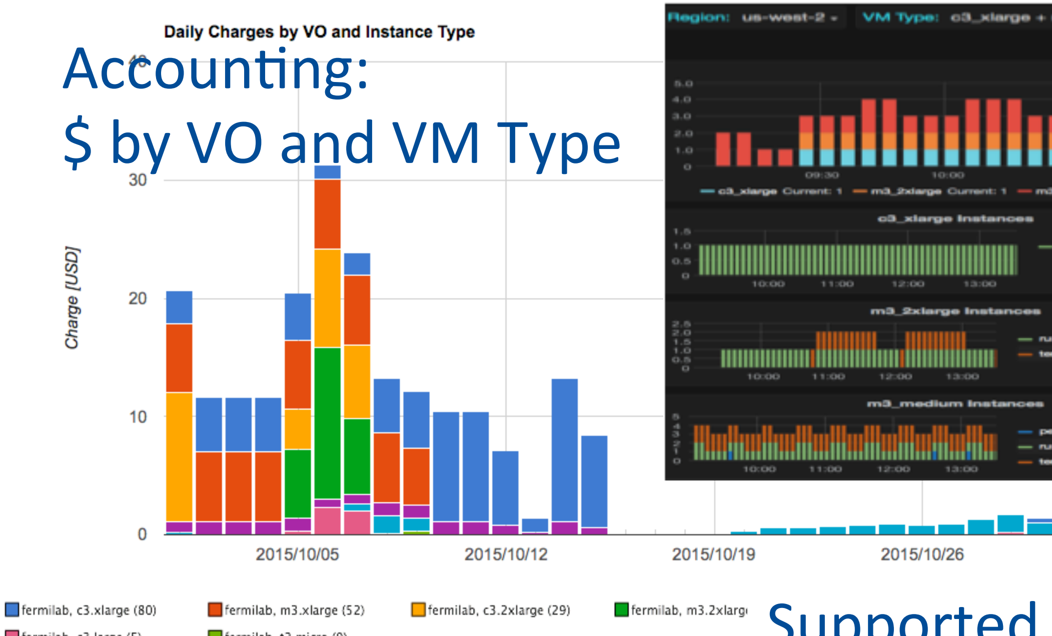

Supported by Shivakumar Vinayagam (IIT)

Who uses HEPCloud?

Some HEPCloud Use Cases

NoVA Processing

Processing the 2014/2015 dataset
16 4-day "campaigns" over one year
Demonstrates stability, availability, cost-effectiveness
Received AWS academic grant

Dark Energy Survey - Gravitational Waves

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)
Modest CPU needs, but want 5-10 hour turnaround
Burst activity driven entirely by physical phenomena (gravitational wave events are transient)
Rapid provisioning to peak

CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) of simulated events in time for Moriond conference
56000 compute cores, steady-state
Demonstrates scalability
Received AWS academic grant

Some HEPCloud Use Cases

NoVA Processing

Processing the 2014/2015 dataset
16 4-day "campaigns" over one year
Demonstrates stability, availability, cost-effectiveness
Received AWS academic grant

Dark Energy Survey - Gravitational Waves

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)
Modest CPU needs, but want 5-10 hour turnaround
Burst activity driven entirely by physical phenomena (gravitational wave events are transient)
Rapid provisioning to peak

CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) of simulated events in time for Moriond conference
58000 compute cores, steady-state
Demonstrates scalability
Received AWS academic grant

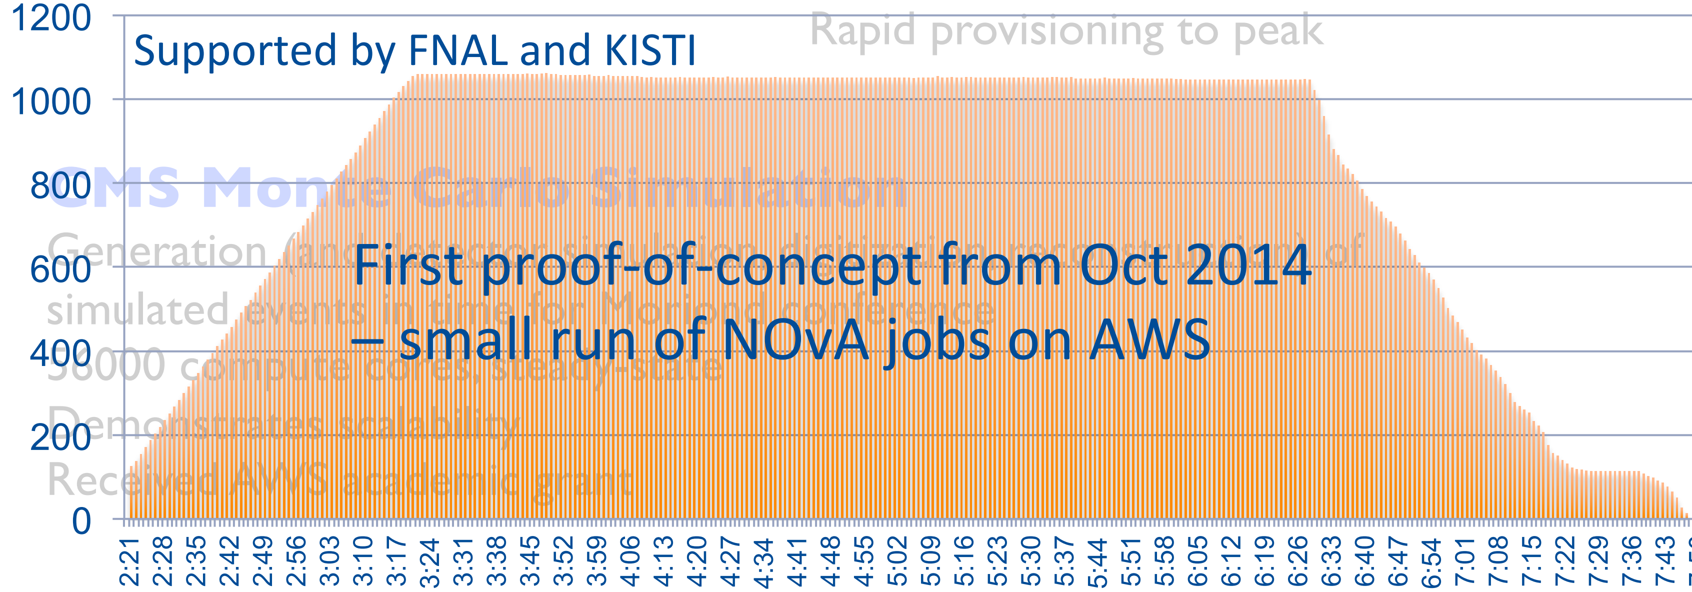

NOvA Use Case – running at 7.3k cores

- Added support for general data-handling tools (SAM, IFDH, F-FTS) for AWS Storage and used them to stage both input datasets and job outputs

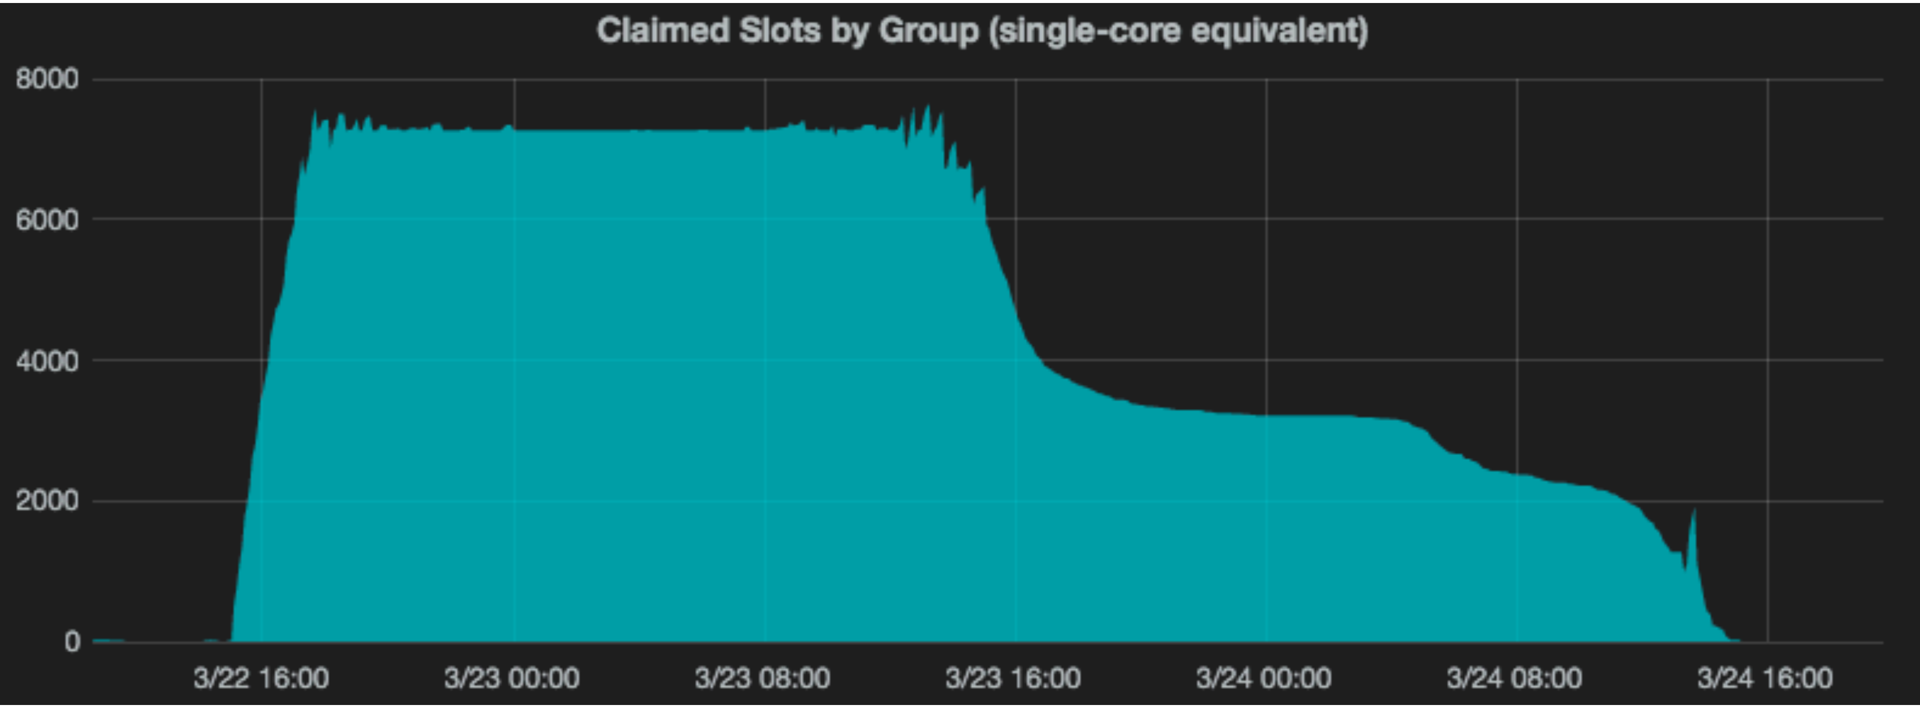

Some HEPCloud Use Cases

NoVA Processing

Processing the 2014/2015 dataset
16 4-day "campaigns" over one year
Demonstrates stability, availability, cost-effectiveness
Received AWS academic grant

Dark Energy Survey - Gravitational Waves

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)
Modest CPU needs, but want 5-10 hour turnaround
Burst activity driven entirely by physical phenomena (gravitational wave events are transient)
Rapid provisioning to peak

CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) of simulated events in time for Moriond conference
56000 compute cores, steady-state
Demonstrates scalability
Received AWS academic grant

Results from the CMS Use Case

- **All CMS requests fulfilled for the “Moriond” conference**
  - **2.9 million jobs, 15.1 million wall hours**
    - 9.5% badput – includes preemption from spot pricing
    - 87% CPU efficiency
  - **518 million events generated**

/DYJetsToLL_M-50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM
/DYJetsToLL_M-10to50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext3-v1/AODSIM
/TTJets_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext1-v1/AODSIM
/WJetsToLNu_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM

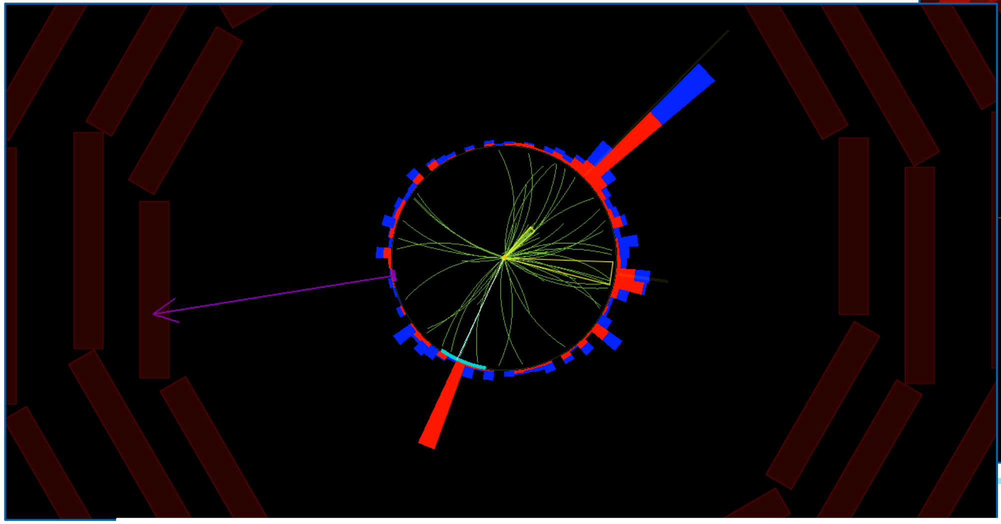

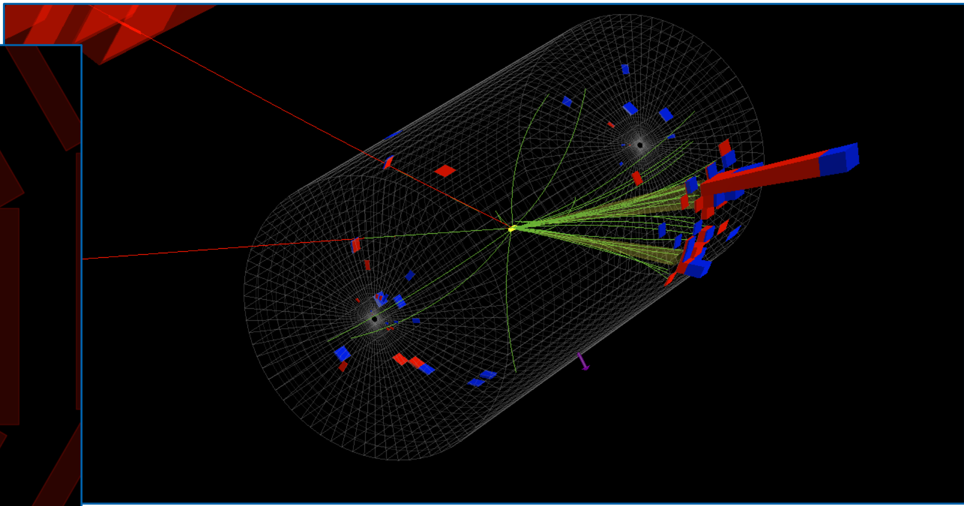

Reaching ~60k slots on AWS with FNAL HEPCloud

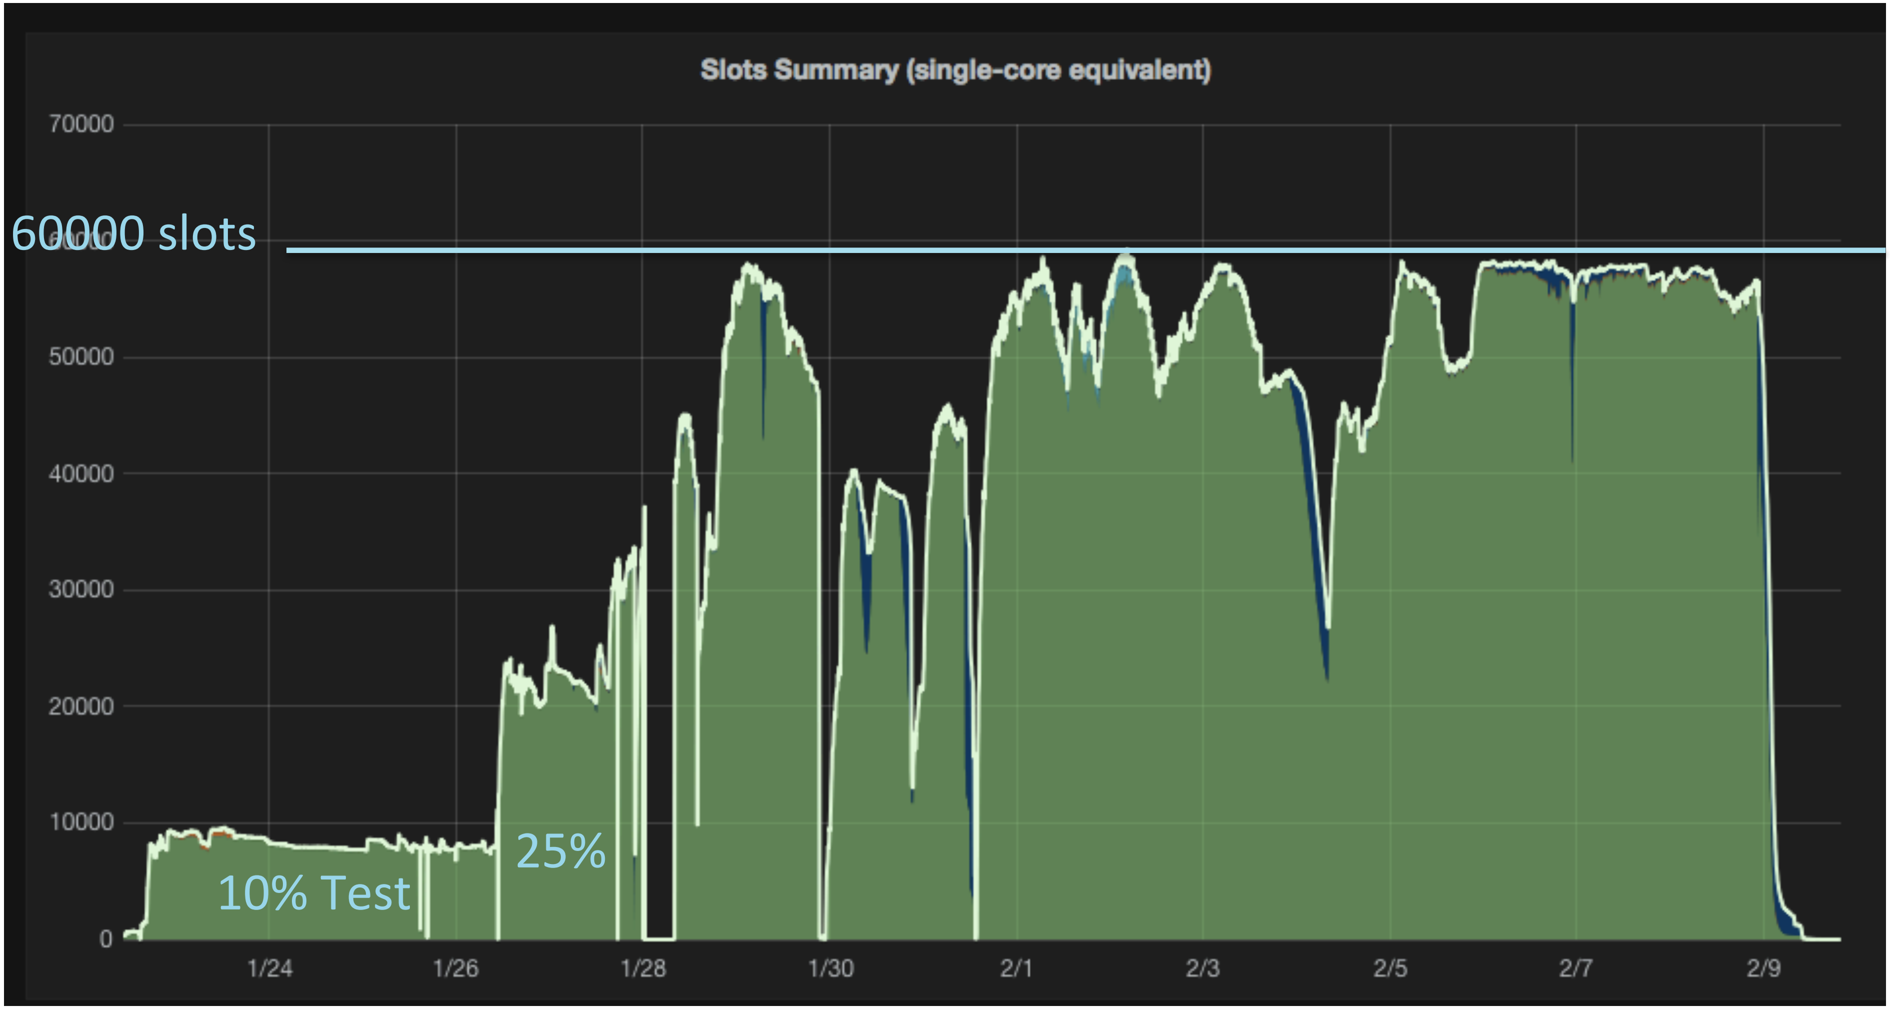

HEPCloud AWS slots by Region/Zone/Type

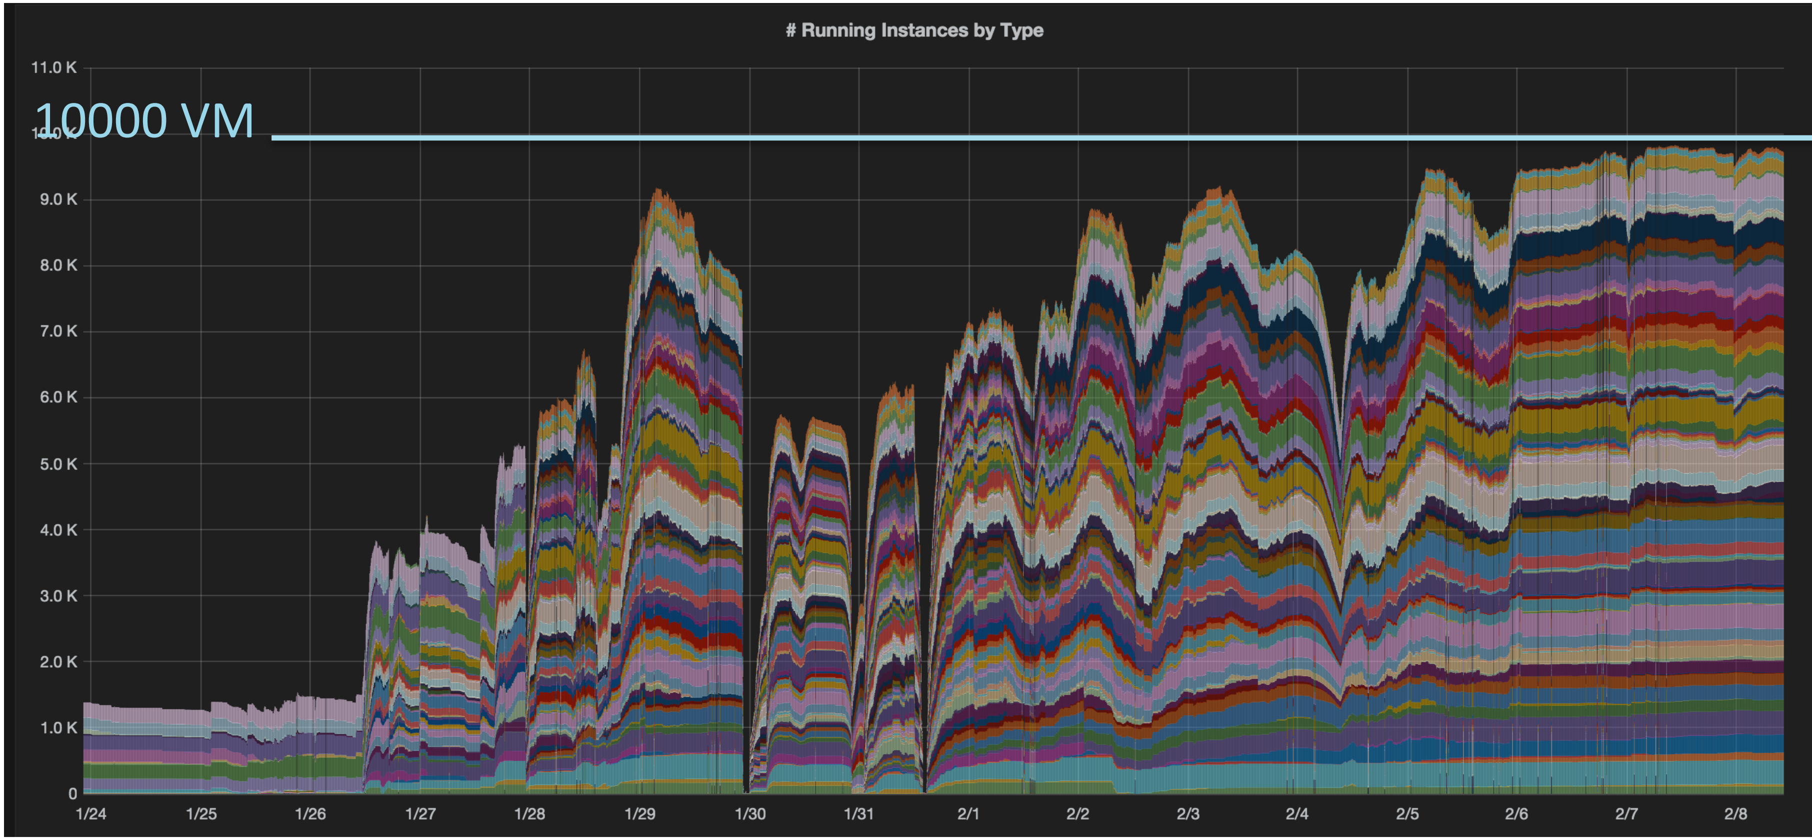

Each color corresponds to a different region+zone+machine type

HEPCloud/AWS: 25% of CMS global capacity

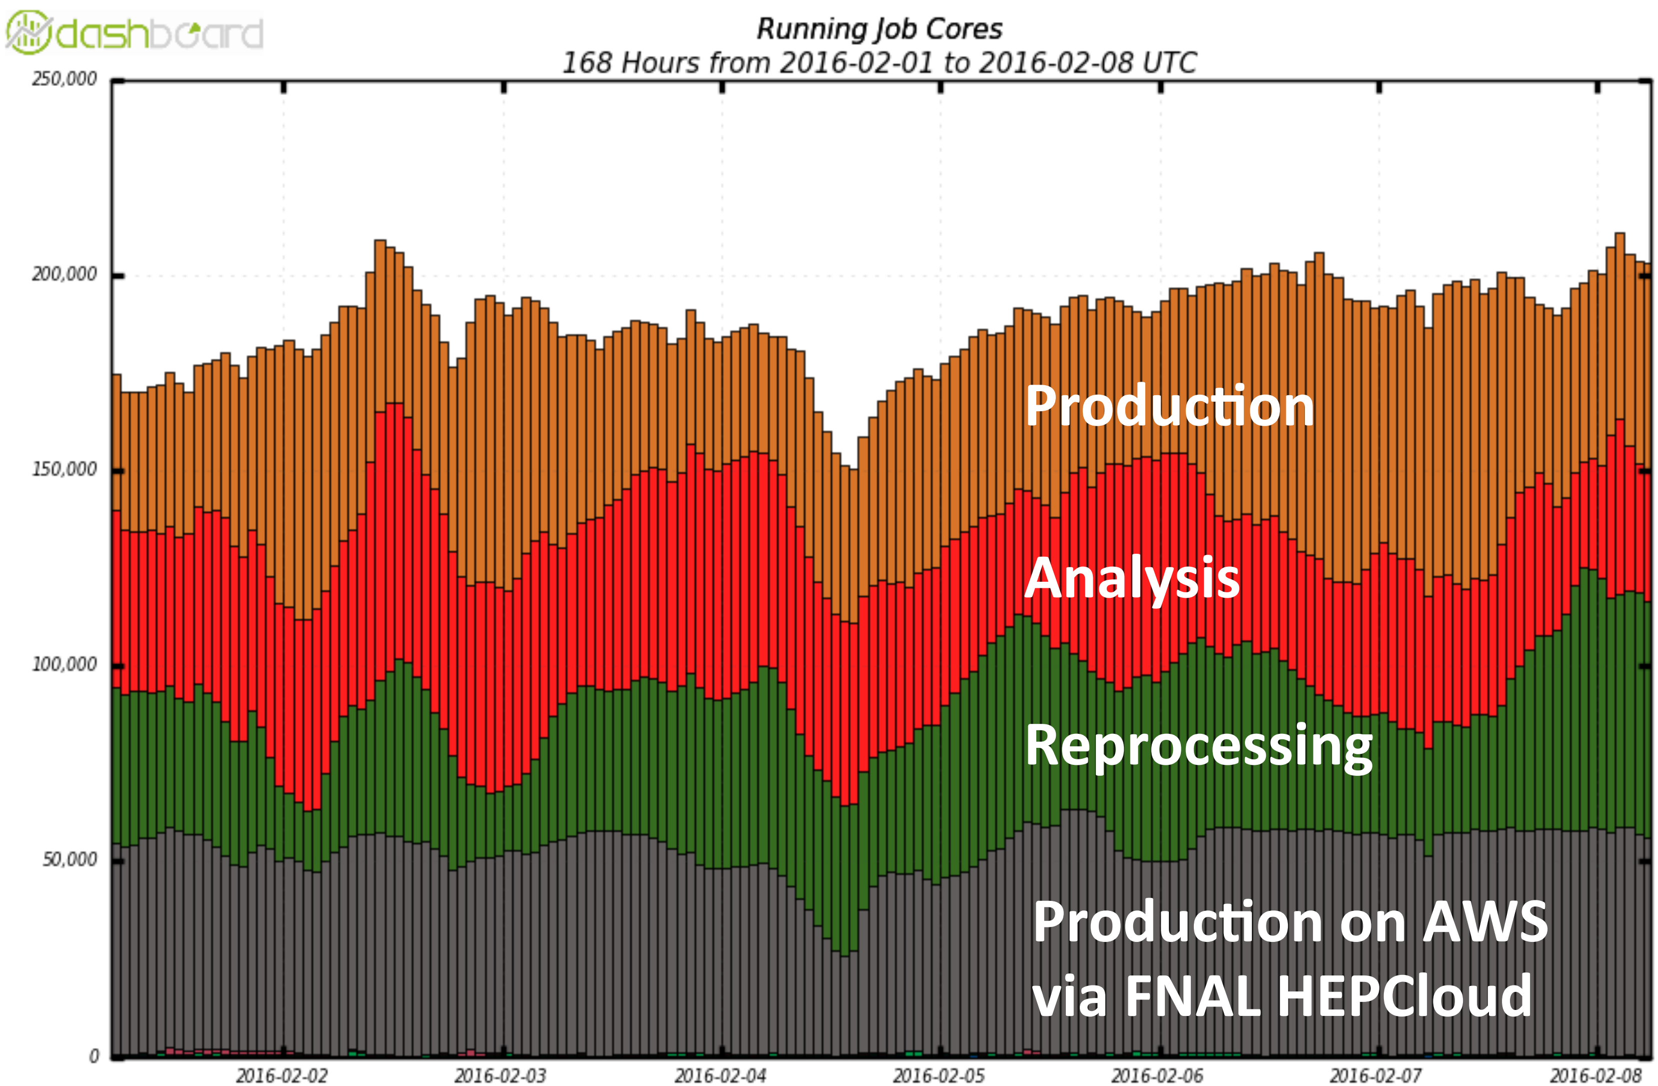

Fermilab HEPCloud compared to global CMS Tier-1

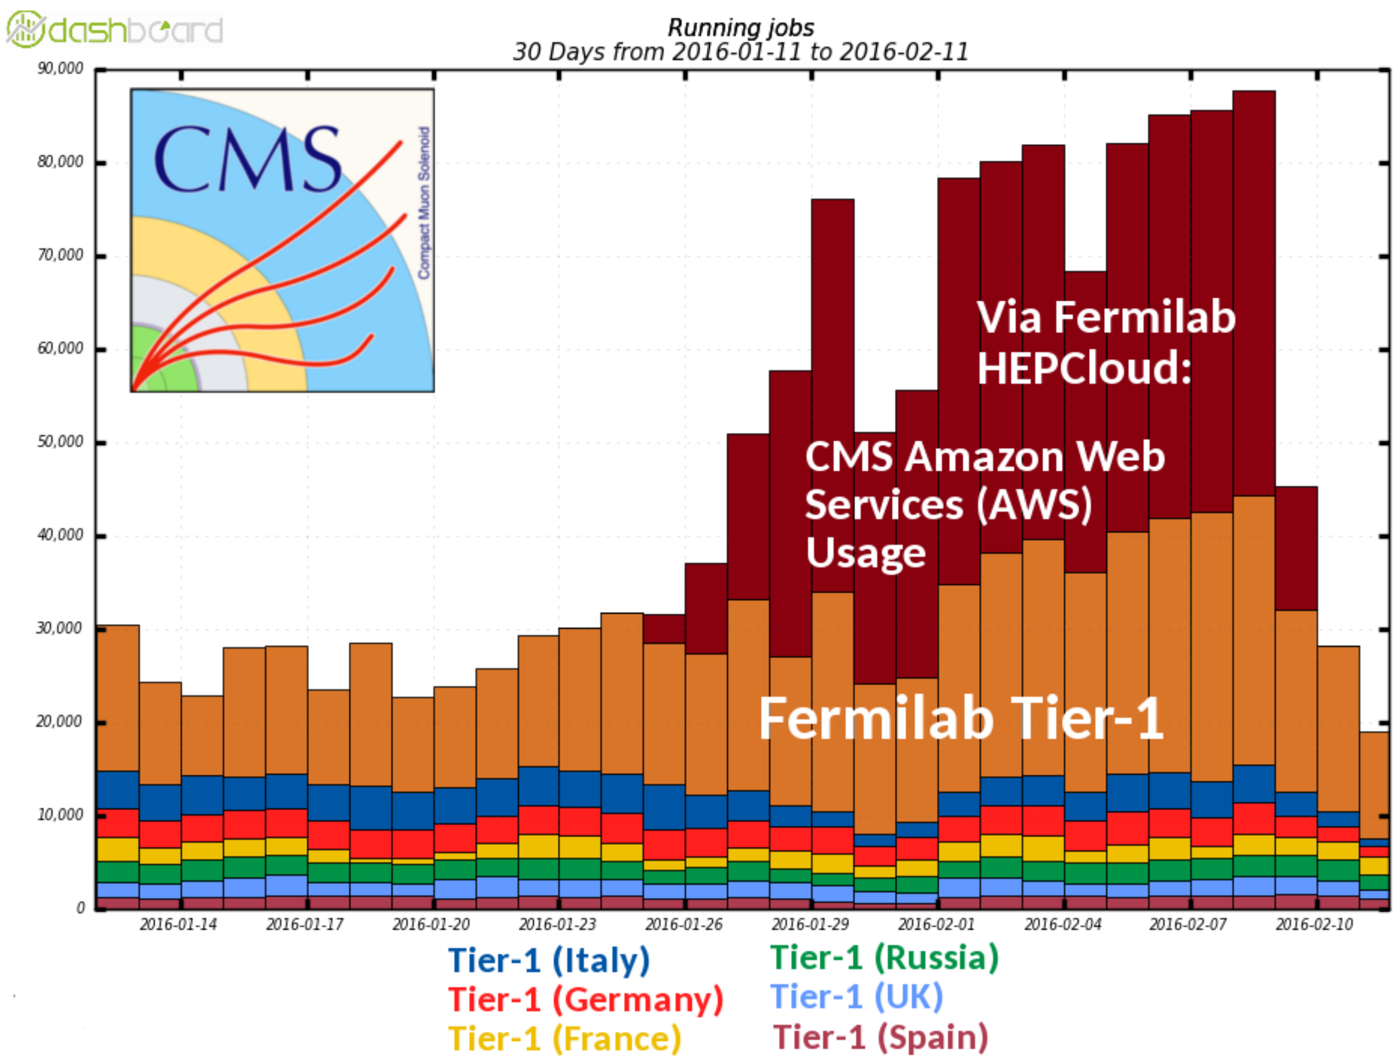

Internship Opportunities

Possible topics for a Fermilab Internship

- Integrate High Performance Computers (HPC) with the HEPCloud Facility. Learn how to execute a simple scientific workflow on a remote HPC site and then work with HEPCloud staff to submit and execute it through the HEPCloud infrastructure.

- Work on functional monitoring of the HEPCloud on-demand-services, both in the AWS cloud and at FNAL. Write monitoring scripts (check_mk) to check the health of the service and integrate the results with the Fermilab site monitoring service. Investigate the forwarding of services logs, such as squid or glideinWMS-pilot services, to a central on-cloud logging server and, from there, to Fermilab.

- Participate in the integration of bulk provisioning systems, such as the condor_annex and AWS Spot Fleet, with the HEPCloud Decision Engine. Add instrumentation to existing monitoring to record the number of instances running and their price. Investigate bulk submission of virtual machines via HTCondor to the Google Compute Engine.

Acknowledgements

- The support from the Computing Sector
- The Fermilab HEPCloud Facility Team and AWS
- The collaboration and contributions from KISTI, in particular Dr. Seo-Young Noh
- The Illinois Institute of Technology (IIT) students and professors Ioan Raicu and Shangping Ren
- The Italian National Institute of Nuclear Physics (INFN) summer student program